

REPORT DOCUMENTATION PAGE				<i>Form Approved</i> OMB No. 0704-0188	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing this collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.					
1. REPORT DATE (DD-MM-YYYY) May 1990		2. REPORT TYPE Conference paper		3. DATES COVERED (From - To)	
4. TITLE AND SUBTITLE See report.				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S) See report.				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) See report.				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) See report.				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION / AVAILABILITY STATEMENT Distribution Statement A - Approved for public release; distribution is unlimited.					
13. SUPPLEMENTARY NOTES Presented at the IEEE 1990 National Aerospace and Electronics Conference (NAECON 1990) held in Dayton, Ohio, on 21-25 May 1990.					
14. ABSTRACT See report.					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT UU	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT Unclassified	b. ABSTRACT Unclassified	c. THIS PAGE Unclassified			19b. TELEPHONE NUMBER (include area code)

TEMPORAL REASONING FOR AIRLIFT SCHEDULING ANALYSIS

Jeffery D. Clay, Capt, USAF
3200 Mission System Division Eglin AFB FL

Michael P. Wellman, Capt, USAF
AI Technology Office
Wright Research and Development Center
Aeronautical Systems Division, WPAFB OH

Charles R. Bisbee, LtCol, USAF
Department of Electrical and Computer Engineering
Air Force Institute of Technology, WPAFB OH

Abstract

Deliberate planning for airlift operations consists of an iterated cycle of resource allocation, simulation, and analysis. The process can be extremely time-consuming for large plans and the final schedule is uncomfortably sensitive to allocation decisions made early in the cycle. These problems can be alleviated by inserting an analysis phase prior to simulation, thereby saving simulation cycles and providing the planner with early feedback about the implications of high-level resource allocation decisions. One useful form of feedback for transportation schedules is an analysis of the temporal constraints among events in the plan. Temporal reasoning provides a general mechanism integrating various types of temporal constraints on airlift events to provide time bounds on execution of the complete plan. For this purpose we developed a general temporal constraint reasoner and a set of mechanisms for deriving temporal information from airlift requirements and partial schedule specifications. Physical limitations of the aircraft and operating facilities as well as the availability of cargo all provide constraints on when certain events may occur. Constraints include the time required to fly from one location to another or the time spent waiting for an aircraft to be loaded. Comparing cargo requirements with airlift capacity over time provides additional constraints. By considering these constraints and their interactions, the planner is in a better position to assess the impact of high-level planning decisions.

1 The Airlift Planning Process

Even the most experienced airlift planners find it difficult to develop an efficient plan for large operations. In a wartime environment, time is critical and days or even hours may determine the difference between success and failure. Developing an effective airlift plan may require several weeks or more. The sheer complexity of the schedule and the number of choices available to the planner contribute significantly to the time required to produce an efficient plan. Because exhaustive enumeration of potential schedules is infeasible, airlift planning follows a hierarchical process. High-level plan descriptions are developed manually and then refined by simulation to produce a final schedule. A seemingly minor choice, made early in the planning process, may make a significant difference in the operational effectiveness of a plan.

Military Airlift Command (MAC) is responsible for the development of airlift plans to support both wartime and

peacetime requirements of the unified and specified commands. There are two types of planning performed at MAC: Crisis Action System (CAS) planning, and deliberate planning. Planning performed during peacetime is called deliberate planning, while CAS planning occurs in response to contingency and crisis situations.

This project (described more fully in [1]) deals only with deliberate planning. The purpose of deliberate planning at MAC is to identify the total movement requirements, to describe them in logistic terms, to simulate the strategic deployment, and to produce a transportation-feasible Operation Plan (OPLAN) [6]. An OPLAN is a plan for the conduct of a single military operation or series of connected operations prepared by the commander of a unified or specified command in response to a requirement established by the Joint Chiefs of Staff.

Military Airlift Command (MAC) uses a program called MACPLAN, developed by MITRE Corp., to aid in developing deliberate airlift plans [7]. MACPLAN provides the human planner with a graphical front-end for specifying high-level resource allocation decisions constraining the ultimately produced airlift schedule. The output of MACPLAN, called a *planset*, consists of a list of operating units and the number and types of aircraft they provide over the course of the airlift operation. Plansets also specify routes between source and destination stations, including designated stations for refueling en route. MACPLAN provides some facilities for analyzing the planset, including an aggregate comparison of airlift capacity and movement requirements over time, and a check of the feasibility of route distances given flying ranges of the chosen aircraft.

Once the planset is generated, it is passed to FLOGEN (FLOW GENERator), which produces a detailed schedule to airlift the requirements within the specifications of the planset. FLOGEN selects aircraft and movement times by sweeping forward in time under the direction of a discrete-event simulator. If FLOGEN fails to produce an acceptable schedule—possibly because the planset is infeasible—the MAC planners must analyze the flow manually, revise the planset, and repeat the process. Given the time required to run FLOGEN and determine the cause of unsuccessful simulations, the overall scheduling activity can be exceedingly long and difficult. Clearly there is much benefit to be gained from automated analysis facilities applied prior to simulation. MACPLAN's facilities provide some of these benefits; this project was designed to extend the range of analysis capabilities available.

2 High-Level Temporal Analysis

One source of constraint on the feasibility of a planset is the temporal bounds on the durations of events in the airlift operation. Analysis of these constraints may provide the planner with feedback about the feasibility and efficiency of a planset without necessarily requiring a detailed schedule. Such feedback may avoid time-consuming applications of FLOGEN and ultimately result in improved airlift schedules.

Our approach is to apply general temporal reasoning methods to temporal information extracted from specifications of transportation requirements and characteristics of aircraft and stations designated in the planset. The transportation requirements dictate the amounts and types of cargo to be moved from one station (onload) to another (offload) within specified periods of time. The time windows are specified in terms of earliest available date, earliest arrival time, and latest arrival time. The station and aircraft databases include such information as the location of the various air bases and the cargo capacity and speed capabilities of the available aircraft types.

Temporal information is in the form of bounds on the time separating distinguished events in the airlift schedule. These can be combined in a network to express temporal constraints among all events of interest. For example, suppose it takes between 10 and 20 minutes to take a shower, between 15 and 20 minutes to eat breakfast and between 25 and 45 minutes to get to work. If you get out of bed at 6:00, you will arrive at work sometime between 6:50 and 7:25. Any additional information constrains the possible event times further. If your shower takes only 10 minutes, your arrival time is now limited to between 6:50 and 7:15. Much work in artificial intelligence has made use of such temporal constraint networks. For a sampling, see [3, 4, 8, 11].

3 Temporal Constraint Propagation

Temporal information is derived from constraint networks via a process of constraint propagation [2, 10]. This process is one of deriving and strengthening constraints among events in the network by combining constraints imposed on other events. The constraint satisfaction paradigm provides a useful framework for describing the constraint propagation method and network representation we employ.

A temporal constraint satisfaction problem (TCSP) as defined by Dechter et al. [4] comprises a set of time points and unary and binary constraints on them expressed in terms of temporal distance. A binary temporal constraint, T_{12} , between the variables X_1 and X_2 indicates the permissible values for the temporal distance $X_2 - X_1$, expressed by disjunctions of inequalities of the form

$$a \leq X_2 - X_1 \leq b.$$

A unary constraint, T_1 , on temporal variable X_1 , dictates the permissible times for the occurrence of X_1 , expressed by disjunctions of inequalities of the form

$$a \leq X_1 \leq b.$$

When all constraints consist of a single disjunct, the TCSP is called *simple*.

A network of temporal constraints can be represented by a directed constraint graph, whose nodes represent temporal variables and whose edges represent constraints on the

temporal distance between the variables. Figure 1 shows a sample temporal constraint network expressing the example described earlier about going to work. X_1 represents waking up, X_2 finishing the shower, X_3 finishing breakfast, and X_4 arriving at work. A solution to this network is a set of values that can be assigned to the edges between the nodes, satisfying all constraints. Assigning the values 15, 18, and 30 to the edges between the nodes in sequential order is one possible solution to this network.

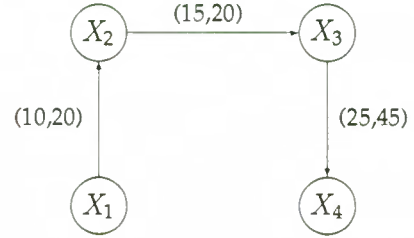


Figure 1: Sample temporal constraint network.

Given a network of temporal constraints, one can answer queries about the possible temporal distances between any pair of variables. If the TCSP is simple, such queries can be answered in time proportional to the cube of the number of nodes [4]. Another approach is to maintain all combinations of relations in the constraint network, which then forms a complete graph. In this approach, queries can be answered in constant time. However, updating the network via constraint propagation upon asserting a new constraint then takes $O(n^3)$ time.

Temporal distance constraints can be composed by simple summation. In the notation of the TCSP framework, $T_{12} = (a, b)$ and $T_{23} = (c, d)$ entails that their *composition* is given by

$$T_{13} = T_{12} \otimes T_{23} = (a + c, b + d). \quad (1)$$

Applying this rule to the relations in Figure 1, we can derive new constraints between X_1 and both X_3 and X_4 . The result is depicted in Figure 2.

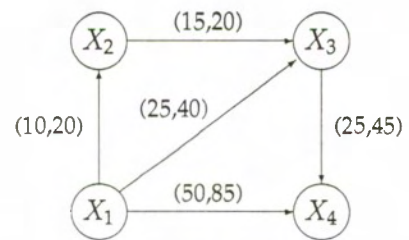


Figure 2: The constraint network after partial propagation.

Even though the complexity is polynomial (cubic), exhaustive propagation of temporal constraints can become infeasible for large constraint networks, which are typical for scheduling problems. One approach to this problem is to cluster the time-points into *reference sets*, limiting propagation to those sets. While the constraint graph for each cluster remains complete, points in distinct clusters may not be directly connected. Queries referring to such separated points are then handled in polynomial time by a search for the shortest path.

Deciding how to cluster the time points into reference sets has a significant impact on the time required to propagate new constraints. Several clustering methods have been proposed in the literature, most based on the temporal relationship among

events. Other methods exploit causality or some other feature represented in the domain knowledge [8]. For application to airlift scheduling, we chose to organize the clusters based on transportation requirements.

4 Airlift Constraint Networks

The first step in our analysis of high-level airlift plans is to translate information in planset form to a temporal constraint network. Constraints from the airlift requirements and capabilities of aircraft and stations involved jointly restrict the temporal course of events in the airlift operation. Analysis of these events provide information about when the operation can be completed, and how much material can be transported during various time-slices of the operation.

To build the constraint network, we first create time points for the significant events in the movement of each airlift requirement. There are five such events for each type of cargo in each requirement: **available-cargo-rx**, **onload-cargo-rx**, **launch-cargo-rx**, **land-cargo-rx**, and **offload-cargo-rx**, where *cargo* is one of the four cargo types—**bulk**, **oversize**, **outsize**, and **pax** (for passenger)—and *x* is replaced by a unique identifier for each requirement. For example, the bulk cargo corresponding to requirement 1 becomes available at time point **available-bulk-r1**. Two special events, **begin-plan** and **end-plan**, serve as boundaries for the airlift operation.

Information in the planset and in the airlift database impose binary constraints between the events in the network. For example, the requirements specify when cargo becomes available, which is encoded as a binary temporal constraint between **begin-plan** and **available-cargo-rx**. Maximum speeds for the various aircraft types combined with the spatial distance of stations place lower bounds on the temporal distance between **launch** and **land** events. Time required for onloading and offloading is similarly dictated by aircraft and station specifications.

Once all of these constraints are asserted into the temporal network, the system can bound the temporal relation between any pair of events. In particular, the time to completion, or *closure*, of the operation is represented by the relation between **begin-plan** and **end-plan**. Similarly, the earliest time that any given requirement can be moved is found by maximizing the earliest time needed for moving its four cargo components.

5 Airlift Capacity Analysis

The constraints described above provide only very loose bounds on the temporal course of events. In effect, the lower bound on the duration of the operation corresponds to the time it would take if aircraft of the fastest type were on hand to move all the cargo as soon as it became available, and stations had the capacity to handle unlimited traffic. We say that the times are *optimistic* because they account for only a subset of the real constraints of the operation. They are a best-case scenario because the constraints represented are valid—it is not possible, according to our airlift model, to complete the operation faster than the constraint network specifies.

We can produce tighter, more realistic bounds on temporal distances among events by incorporating constraints accounting for the limitations of aircraft and station resources. Enforcing such constraints, however, is difficult in the context of a high-level plan, as there has not yet been a commitment

to the allocation of resources to individual requirements. The problem is that resource limitations constrain our ability to satisfy combinations of requirements, although any individual requirement could be satisfied in isolation. Therefore, an encoding of these constraints in the temporal constraint network would require substantial use of disjunction, which would mean that the problem is no longer a simple TCSP, and thus computationally intractable [4].

One approach is to perform an analysis of airlift capacity in the aggregate, without necessarily linking the result to the individual movement events of the constraint network. This is the approach taken by the capacity analysis module of MACPLAN, which compares capacity with airlift requirements over time, where each is measured in terms of gross cargo ton-miles and passenger-miles. MACPLAN computes capacity for a given day as a function of the aircraft available, considering the size, speed, and utilization rates of the various aircraft types. The program compares this capacity with the average requirements over time, computing a profile of the backlog of material to be transported throughout the operation. Closure occurs when the backlog reaches zero.

The capacity analysis of MACPLAN is also optimistic because the backlog computations in effect assume that the available aircraft can be used at full capacity. In the simulation produced by FLOGEN, this will not be the case, as the geographical and temporal distribution of requirements invariably necessitate flights with less-than-full loads. Nevertheless, the analysis provides an upper bound on capacity, and therefore a lower bound on backlog and closure. These can be asserted in the constraint network in conjunction with the information discussed in Section 4.

(Note that for purposes of discussion here, our gold standard for optimism and the validity of temporal bounds is the range of possible simulations producible by FLOGEN. The actual airlift operation may deviate from the analysis to the extent that the reality deviates from the data in the given airlift model. For example, in a real airlift, the operating wing may override fleet utilization rates based on specific, locally available information on aircraft, crews, and other factors.)

The question remains whether we can account for the fragmentation of requirements and underutilization of capacity to produce more precise temporal bounds on the airlift events, without first producing a detailed schedule. Our approach is to analyze capacity based on a partial matching of resources to requirements designed to conservatively account for underutilization. That is, we expect that the transportation capacity computed from this partial description will be greater than for any feasible complete schedule for the given requirements. If this relation is guaranteed (that is, the capacity analysis is optimistic), then the method produces valid lower bounds on closure and backlog over time.

We have experimented with a variety of procedures for producing the approximate capacity analyses. None are guaranteed to be optimistic, so it is possible that actual schedules could airlift the requirements faster. The general method is to match requirements to aircraft and compute the total cargo and passengers moved under that assignment. The accuracy and degree of optimism of these estimates depends on two factors:

1. The assumptions restricting the range of possible assignments.
2. The likelihood that this assignment has capacity as great as any other meeting the assumptions.

The estimate becomes more accurate as the underlying assumptions reflect a truer set of constraints on the operation. Its optimism is more assured as the assumptions are less restrictive and the assignment is more selective. Restricting the assumptions produces capacities less likely to be optimistic and therefore provides stronger temporal lower bounds at the expense of confidence in their validity.

The first assumption implicit in our assignment algorithm is that an aircraft can only be allocated to a single requirement on a given day. This does not reflect an absolute constraint, so this assumption is a potential source of non-optimism. Nevertheless, we chose to adopt it because it appears to hold generally, it significantly simplifies the assignment process, and it provides a direct means to account for underutilization of aircraft capacity due to requirement fragmentation. A compensating factor is that the assignments do not reflect the constraint on locations of aircraft from day to day. This admits some non-feasible assignments, thus improving the likelihood that the selected one will provide a capacity at least as great as any that are feasible.

The assignment process operates day-by-day with a list of requirements and aircraft available. Each type of cargo is considered separately and each is allotted aircraft from the total number available for that day. If two C-5s and two 747s are available, for instance, all four aircraft are considered when moving the bulk cargo and all are reconsidered when moving the passenger, oversize, and outsize cargo. The convention of not forcing attrition of the identified aircraft after moving one type of cargo also contributes to a better-than-best-case (optimistic) assignment.

The assignment for a given day is based on a greedy allocation of requirements to the available aircraft. Each aircraft is assigned the greatest amount of cargo remaining in a single requirement, up to its full capacity. Although the greedy assignment is not necessarily optimal, it serves as a reasonable and simple approximation.

A variety of assignments can be produced by varying the conventions for deeming requirements eligible for transport. One method we applied matched aircraft for a given day against a cumulative list of requirements for all previous days. A second method deleted a requirement from the list if it was selected for movement on a previous day. The former method generally produces infeasible assignments where requirements are moved more than once, while the latter is subject to suboptimality due to myopic selection of requirements. Thus, the first method is more optimistic, while the second is likely to produce stronger bounds.

6 Results

As expected, the backlog and temporal lower bounds increased as we relaxed the optimism criterion. Since MACPLAN's capacity analysis is optimistic, both of our methods described above produce greater backlogs. While these bounds are not guaranteed, the conservatism of other assumptions of the method provide reason to believe that these greater backlogs will be more accurate predictors of the capacity profile of the detailed produced via simulation. By varying the assumptions underlying the analysis procedure one can trade off confidence in the validity of the bounds for expected accuracy of the capacity figures.

Further analysis and empirical study are required to determine the precise form of the tradeoff and to choose the

capacity estimation techniques offering the most appropriate feedback to MAC planners. Once selected, it would be easy to incorporate these techniques within the existing MACPLAN framework.

7 Limitations and Extensions

The temporal constraint network and capacity analysis system provide a framework for temporal analysis of plansets, or high-level airlift plans. The capability provided by existing facilities, however, is quite limited. Further work on extensions to the system will be required to fulfill the potential for meaningful support to airlift planners.

The temporal reasoning subsystem can be improved in several ways. First, we could make better use of the facility for representing disjunctive constraints. This would provide a direct way of expressing incomplete information without requiring unjustified commitments to the allocation of resources. This expressive power is provided in our implementation, though the reasoning capabilities for this type of constraint are limited. Dechter et al. [4] discuss a variety of strategies for disjunctive temporal constraint satisfaction.

A second area of improvement is in the clustering of events into reference sets. The complexity of constraint propagation is most sensitive to the size of reference sets. Our clustering scheme, based on requirements and cargo types, has been satisfactory for our work to date. However, alternate strategies for clustering will probably be required once constraints on events spanning different requirements become more commonplace.

Our final recommendation for the temporal constraint system is a facility for reason maintenance [5]. Given rapidly changing information or a range of choices that a planner wishes to explore, a facility for retracting previously asserted constraints can have significant value.

The techniques for capacity analysis can also be improved. Potential enhancements include the cargo/aircraft matching algorithm and the aircraft and station availability calculations. The greedy allocation of requirements to aircraft may not be optimal, which compromises the optimism of the final result. A variety of approaches offer potential remedies to the problem, including heuristic search, inter-temporal matching, and bounded approximations of the optimal solution. We would also pursue the possibility of abandoning the match altogether in favor of other methods for estimating aggregate airlift capacity.

One source of constraint not considered by the capacity analysis system is the limitations in station resources for parking, refueling and other airlift functions. To exploit these constraints we could conduct an explicit search for bottlenecks in a station-centered view of the scheduling problem [9]. This view may also help us account for geographic distance in assessing the day-to-day availability of aircraft for the airlift operation.

The utilization rate of each aircraft type determines how many hours, on average, planes of that type are available to carry cargo on each day. Analyzing capacity by utilization rate is a relatively crude method, as it applies the constraint to the fleet as a whole rather than individual aircraft. Accounting for the fundamental constraints affecting an aircraft's availability (for example, maintenance and crew scheduling) may yield a more precise profile of airlift resources.

8 Conclusion

Enhanced automated analysis at any level would be a boon to deliberate airlift planners. Analysis of high-level plansets is particularly valuable because it provides useful feedback early in the planning process. Our work on the problem to date suggests that some forms of capacity analysis based on temporal constraints can be achieved, and that such analyses can be integrated straightforwardly with existing planning support systems. Moreover, a variety of sources of constraint remain untapped. Further work on methods to account for some of these constraints would produce more precise analyses and thereby contribute to a more efficient and competent planning process.

Acknowledgment

We thank Joe Besselman of Electronic Systems Division and Murray Daniels, Kirsten Kissmeyer, and Anne Tallant of MITRE for their assistance in working with MACPLAN.

References

- [1] Jeffery Dean Clay. Temporal constraint propagation for airlift planning analysis. Technical Report AFIT/GCE/ENG/89D-1, Air Force Institute of Technology, Wright-Patterson AFB, OH 45433, December 1989.
- [2] Ernest Davis. Constraint propagation with interval labels. *Artificial Intelligence*, 32(3):281–331, 1987.
- [3] Thomas L. Dean and Drew V. McDermott. Temporal data base management. *Artificial Intelligence*, 32:1–55, 1987.
- [4] Rina Dechter, Itay Meiri, and Judea Pearl. Temporal constraint networks. In *First International Conference on Principles of Knowledge Representation and Reasoning*, pages 83–93, 1989.
- [5] Jon Doyle. A truth maintenance system. *Artificial Intelligence*, 12(2):231–272, 1979.
- [6] HQ MAC/DO-ADANS. ADANS functional description. Technical report, Headquarters, Military Airlift Command, Scott AFB, IL, November 1988.
- [7] Kirsten Y. Kissmeyer and Anne M. Tallant. MACPLAN, a mixed-initiative approach to airlift planning. In Herbert Schorr and Alain Rappaport, editors, *Innovative Applications of Artificial Intelligence*. MIT Press, 1989.
- [8] Isaac S. Kohane. Temporal reasoning in medical expert systems. In R. Salamon, B. Blum, and M. Jørgensen, editors, *MEDINFO 86: Proceedings of the Fifth Conference on Medical Informatics*, pages 170–174, Washington, October 1986. North-Holland.
- [9] Peng Si Ow, Stephen F. Smith, and Alfred Thiriez. Reactive plan revision. In *Proceedings of the National Conference on Artificial Intelligence*, pages 77–82. AAAI, 1988.
- [10] Gerald Jay Sussman and Guy Lewis Steele Jr. CONSTRAINTS—A language for expressing almost-hierarchical descriptions. *Artificial Intelligence*, 14:1–39, 1980.
- [11] Marc B. Vilain. A system for reasoning about time. In *Proceedings of the National Conference on Artificial Intelligence*, pages 197–201. AAAI, 1982.